arXiv:1407.7614v1 [stat.ME] 29 Jul 2014

Confidence Areas for Fixed-Effects PCA

Julie Josse
Applied Mathematics Department, Agrocampus Ouest, Rennes, France
and
Stefan Wager
Department of Statistics, Stanford University, California, USA
and
François Husson
Applied Mathematics Department, Agrocampus Ouest, Rennes, France

July 30, 2014

Abstract

PCA is often used to visualize data when the rows and the columns are both of interest. In such a setting there is a lack of inferential methods on the PCA output. We study the asymptotic variance of a fixed-effects model for PCA, and propose several approaches to assessing the variability of PCA estimates: a method based on a parametric bootstrap, a new cell-wise jackknife, as well as a computationally cheaper approximation to the jackknife. We visualize the confidence regions by Procrustes rotation. Using a simulation study, we compare the proposed methods and highlight the strengths and drawbacks of each method as we vary the number of rows, the number of columns, and the strength of the relationships between variables.

Keywords: confidence ellipses, fixed-effects model, jackknife, parametric bootstrap, PCA, variability

1 Introduction

Principal component analysis (PCA) is a well-established dimensionality reduction method used to explore and to visualize data sets. PCA is often performed on data sets that can be characterized as "population data sets," where both the rows and the columns of the data matrix are fixed. Even in this case of population data sets, we may assume that the observed data has been corrupted by

random noise, and so we may want to to perform statistical inference of the PCA output. In this paper, we develop methods that let us understand the impact of noise on PCA, and to visualize this variability on PCA graphical outputs using confidence ellipses.

To give a concrete example, suppose that our dataset consists of average temperatures for the major cities in a country (rows) across the twelve months (columns); our goal is to describe the data using PCA and, e.g., to visualize which cities have similar temperature profiles and to study the correlation between the temperatures. In this example, neither the cities nor the months can be treated as random; however, it is natural to assume that there is some noise in our long-term average temperature measurements. The confidence areas developed in this paper enhance the interpretability of the PCA output. For example, they would enable us to see which cities are significantly different with respect to their climates and which cities are not.

The appropriate way of doing inference is heavily dependent on the noise model for the data matrix $\mathbf{X}$. Classical inference approaches for PCA often assume that the rows of the data matrix $\mathbf{X}$ form a random sample from a larger population. In this case, one of the major approaches for assessing the variability of PCA is to use non-parametric bootstrap procedures (Daudin et al, 1988; Holmes, 1989; Chateau and Lebart, 1996; Timmerman et al, 2007; Badamoradi et al, 2013), which involves repeatedly re-sampling the rows of $\mathbf{X}$ with replacement. Then, two strategies were suggested for visualizing the variability on the PCA graphical outputs: the "partial bootstrap" and the "total bootstrap" (Greenacre, 1984; Chateau and Lebart, 1996). The partial bootstrap consists in projecting the B replications onto the initial subspace (obtained by the PCA applied on the initial data matrix), while the total bootstrap consists in performing a PCA on each bootstrap replication. However, in order to compare the results provided by each PCA, it is necessary to minimize the "nuisance variations" (Milan, 1995), meaning the possible reflections or rotations between different configurations. To this end, it is possible to use Procrustes rotations (Gower and Dijksterhuis, 2004) to fit the B bootstrap PCA configurations toward the initial one. Rotation problems are avoided with the partial bootstrap but Milan (1995) showed that this procedure underestimates the variability of the parameters. Such a framework enables us to study the variability of the columns coordinates (or the loadings) and is well suited, for example, to situations where PCA is applied to survey data sets in which each row represents an anonymous survey participant.

We focus on an alternative noise model for the data called the "fixed-effects" noise model where both the rows and the columns of $\mathbf{X}$ are non-random, and randomness comes only from measurement error. It is in agreement with the cities example where the rows of $\mathbf{X}$ represent specific subjects whom we want to study, and the columns of $\mathbf{X}$ represent different features of the subjects. In this case, we cannot reasonably resample the rows of $\mathbf{X}$, because the subjects we chose to study are non-random. Many examples of such data can be found in many fields. We will analyze in Section 6.2 a data set from sensory analysis where the rows are food products (different wines) and the columns sensory descriptors (sweetness, bitterness, etc.).

In this paper, we develop and study methods for producing confidence areas for PCA in such fixed-effects models. We begin in Section 2 by defining a fixed-dimensional noise model for PCA. We then use this noise model to build confidence areas using various methods. In Section 3, we define asymptotic confidence regions for PCA areas using classical theory for non-linear regression. The variability of models related to the fixed-effects PCA has been studied by Denis and Gower (1994) and Denis and Pazman (1999), and we use their results to build these asymptotic confidence areas. We then present alternative methods based on both a parametric bootstrap (Section 4.1) and a cell-wise jackknife (Section 4.2). As the jackknife-based method can be computationally intensive, we also propose an approximation to the jackknife based on the "leave-out-one" lemma of Craven and Wahba (1979). In order to use these confidence areas in practice, we need a way to translate them into visualizations. In Section 5, we discuss Procrustes rotations that tackle problems related to the translation and rotation of different PCA configurations. Finally, in Section 6, we assess the coverage properties of the proposed methods using an extensive simulation study and illustrate our methods on real datasets.

2 Two Perspectives on Principal Component Analysis

As discussed above, in order to do inference for fixed-effects PCA we first need to define a fixed-dimensional noise model. We propose such a model in Section 2.2. Before doing so, however, we first review the geometric argument that is often used to motivate the PCA procedure.

2.1 The Geometric Point of View

PCA provides a low-dimensional space which minimizes the distances between points and their projections onto this subspace. More precisely, let $\mathbf{X}_{n\times p}$ be a (centered) matrix, and let $\|\mathbf{A}\| = \sqrt{\mathbf{tr}(\mathbf{A}\mathbf{A}')}$ be the Hilbert-Schmidt norm. Then, PCA corresponds to finding a matrix $\hat{\mathbf{X}}_{n\times p}$ of rank S that provides the best approximation to the matrix $\mathbf{X}$ under least squares norm, i.e., which minimizes $\|\mathbf{X} - \hat{\mathbf{X}}\|^2$. The solution is given by the singular value decomposition (SVD) of $\mathbf{X}$:

$$\hat{\mathbf{X}}^{(S)} = \mathbf{U}_{n\times S}\mathbf{\Lambda}^{\frac{1}{2}}_{S\times S}\mathbf{V}'_{p\times S},$$

where $\mathbf{U}_{n\times S}$ and $\mathbf{V}_{p\times S}$ are the matrices of left and right singular vectors and $\mathbf{\Lambda}_{S\times S}$ is the diagonal matrix with the associated singular values. The number of components S is considered here to be known. The matrix $\mathbf{V}$ is often known as the loadings matrix, the coefficients matrix or the principal axes matrix; the matrix $\mathbf{F} = \mathbf{U}_{n\times S}\mathbf{\Lambda}^{\frac{1}{2}}_{S\times S}$ is known as the scores matrix, the principal components matrix or as the rows coordinates matrix.

2.2 The Model-Based Point of View

PCA can also be presented with a model where the data are generated as a structure corrupted by noise:

$$\mathbf{X} = \tilde{\mathbf{X}} + \boldsymbol{\varepsilon},\ \varepsilon_{ij} \sim \mathcal{N}(0, \sigma^2) \tag{1}$$

The structure can be considered to be random or fixed. Caussinus (1986) suggested a model where the data are generated as a fixed structure having a low rank representation in S dimensions corrupted by noise:

$$x_{ij} = \sum_{s=1}^{S} \sqrt{d_s} q_{is} r_{js} + \varepsilon_{ij},\ \ \varepsilon_{ij} \sim \mathcal{N}(0, \sigma^2) \tag{2}$$

In this model the rows have different expectations and the randomness is only due to the error term. Thus, it is more in agreement with cases where both the rows and the columns of the data matrix are of equal interest. The maximum likelihood estimates correspond to the least squares solution, i.e., to the principal axes and components ($\hat{x}^{(S)}_{ij} = \sum_{s=1}^{S} \sqrt{\lambda_s}\mathbf{u}_{is}\mathbf{v}_{js}$). This model serves as a basis to define procedures to study the variability in PCA.

Note that the noise model (1) can also be paired with a random structure for $\mathbf{X}$. This leads to procedures like probabilistic PCA (Tipping and Bishop, 1999), or the more general Factor

Analysis model (Bartholomew, 1987). Such random effects models are more closely in agreement with situations where the non-parametric bootstrap procedure, described in the introduction, is appropriate.

3 Asymptotic Theory for Fixed-Effects PCA

In the inferential framework associated with (2), the number of parameters that need to be estimated grows proportionally with the size of $\mathbf{X}$. Thus, the asymptotic regime is only reached by taking the variance of the noise ε down to 0. In our example presented in Section 6.2, the entries x_{ij} are obtained by averaging the responses from multiple consumer surveys. Thus, we will naturally reach $\sigma^2 \to 0$ asymptotics as we survey more and more people.

Asymptotic confidence areas in PCA can be obtained using results originally developed for the two-factor analysis of variance problem. We build on the work of Denis and Gower (1994) and Denis and Pazman (1999), who studied non-linear regression models with constraints, with a focus on bi-additive models of the form

$$y_{ij} = \mu + \alpha_i + \beta_j + \sum_{s=1}^{S} \gamma_{is}\delta_{js} + \varepsilon_{ij} \text{ with } \varepsilon_{ij} \sim \mathcal{N}(0, \sigma^2). \tag{3}$$

Here, y_{ij} is the response for the $(i,\ j)$ category pair, μ is the grand mean, $(\alpha_i)_{i=1,\ldots,I}$ and $(\beta_j)_{j=1,\ldots,J}$ correspond to the main effect parameters, and

$$\left(\sum_{s=1}^{S} \gamma_{is}\delta_{js}\right)_{i=1,\ \ldots,\ I;\ j=1,\ \ldots,\ J}$$

models the interaction. Such models are useful for studying the interaction between two factors when no replication is available, and have been successively applied in agronomic fields such as crop-science (Gauch, 1988; Gauch and Zobel, 1996; Cornelius et al, 1996). Denis and Gower (1994, 1996) derived the asymptotic variance of maximum likelihood estimators in the model (3). Denis and Pazman (1999) then used these results to compute the bias of the parameters of model (3) and showed that even if the estimators of the parameters are biased, the approximate bias of the response estimator is null. Similar results have more recently been rediscovered by Papadopoulo and Lourakis (2000). These results let us directly derive the asymptotic distribution of the PCA estimator in model (2). Indeed, from a computational point of view, model (3) is similar to the

PCA model (2), the main difference being that the linear part only includes the grand mean and column main effect in PCA. Thus, we find that the PCA estimator is asymptotically unbiased $\mathbb{E}(\hat{x}_{ij}^{(S)}) = \tilde{x}_{ij}$ and that the variance of the first order asymptotic approximation is

$$\begin{aligned} &\mathbb{V}\left(\hat{x}_{ij}^{(S)}\right) = \sigma^2 \mathbf{P}_{ij,ij}, \text{ where} \\ &\mathbf{P} = \left(\mathbb{I}_p \otimes \frac{1}{n}\mathbb{1}\mathbb{1}'\right) + \left(\mathbf{P}_{\mathbf{V}}' \otimes \left(\mathbb{I}_n - \frac{1}{n}\mathbb{1}\mathbb{1}'\right)\right) + (\mathbb{I}_p \otimes \mathbf{P}_{\mathbf{U}}) - \left(\mathbf{P}_{\mathbf{V}}' \otimes \mathbf{P}_{\mathbf{U}}\right). \end{aligned} \tag{4}$$

Here, $\mathbb{I}_p$ (respectively $\mathbb{I}_n$) is the identity matrix of order p (resp. of order n), $\mathbb{1}$ is the n-vector of 1's, $\otimes$ is the Kronecker product, and

$$\mathbf{P}_{\mathbf{U}} = \mathbf{U}(\mathbf{U}'\mathbf{U})^{-1}\mathbf{U}' \text{ and } \mathbf{P}_{\mathbf{V}} = \mathbf{V}(\mathbf{V}'\mathbf{V})^{-1}\mathbf{V}'.$$

$\mathbf{P}_{\mathbf{U}}$ and $\mathbf{P}_{\mathbf{V}}$ are the two orthogonal projection matrices involved in PCA, representing projections onto the spaces spanned by $\mathbf{U}$ and $\mathbf{V}$ respectively. These results directly lead to Gaussian asymptotic confidence regions. To compute these regions in practice, we draw matrices

$$\left(\hat{\mathbf{X}}^{(S)1}, ..., \hat{\mathbf{X}}^{(S)\star}\right)$$

from a Gaussian distribution with expectation $\hat{\mathbf{X}}$ and variance given by equation (4). We will discuss in Section 5 how to use these matrices to draw confidence areas around the row and column points in PCA graphical outputs.

The expression (4) can also be derived by writing PCA as a smoothing operator. More precisely, it can be shown (Candes and Tao, 2009; Josse and Husson, 2011) that PCA can be written as follows:

$$\text{vec}(\hat{\mathbf{X}}^{(S)}) \quad = \quad \mathbf{P}\text{vec}(\mathbf{X}), \tag{5}$$

where vec is the vectorization operator, i.e., $\text{vec}(\hat{\mathbf{X}}^{(S)})$ is a vector of size np with the columns of $\hat{\mathbf{X}}^{(S)}$ stacked below each other. Here, the matrix $\mathbf{P}$ from (4) acts as an orthogonal projection matrix whose elements depend on the data $\mathbf{X}$. Consequently, PCA can be seen as a "non-linear" model where $\mathbf{P}$ represents the projection onto the tangent space of the expectation surface. Thus, as in classical non-linear regression models, we can obtain asymptotic forms by studying linear approximations.

Confidence areas based on linear approximations are expected to be valid as long as the non-linearity is small. Extending prior work by Bates and Watt (1980) and Pazman (2002), Pazman

and Denis (2002) define an intrinsic measure K_{int} of non-linearity for biadditive models (3):

$$K_{int} = \frac{1}{\lambda_S^{1/2}}, \text{ with } \lambda_s^{1/2} = \sum_i \gamma_{is}^2 = \sum_j \delta_{js}^2.$$

A small value of σK_{int} implies better accuracy of asymptotic confidence regions (Bates and Watt, 1980). Thus, we should expect that the asymptotic confidence areas for PCA perform best when the signal-to-noise ratio (SNR) is high. In addition, the validity of the asymptotics is all the more reliable when n and p are large. We assess these results in a simulation study in Section 6.

4 Confidence Areas for Fixed-Effects PCA

In this section, we propose alternatives to the asymptotic confidence areas derived above. Our methods are based both on a parametric bootstrap and the jackknife. In small samples, these non-parametric confidence areas can often out-perform their asymptotic counterpart.

4.1 The Parametric Bootstrap

Bootstrap methods are often used to obtain confidence areas when the underlying model is too difficult to study analytically or before we enter the asymptotic regime. For random-effect designs where the non-parametric bootstrap procedure described in the introduction is appropriate, Timmerman et al (2007) showed that the bootstrap is more flexible and gives better confidence areas in the sense of the coverage than the corresponding asymptotic rule.

In the case of a fixed-effects model (2), however, we cannot reasonably resample the rows of $\mathbf{X}$ as the structure of $\mathbf{X}$ is non-random. But we can still define a bootstrap procedure by regenerating residuals. This approach is called a parametric bootstrap (Efron and Tibshirani, 1994) as it requires an explicit model. It allows us to study the variability of the parameters due to the noise ε in (2), and to understand how the estimated parameters might have behaved with a different realization of the noise. This method allows us to study parameters related to both the scores and the loadings, as opposed to the non-parametric bootstrap which only lets us look at the loadings.

We use the following parametric bootstrap algorithm:

1. Perform the PCA on $\mathbf{X}$ to estimate the parameters $\mathbf{U}_{n\times S}$, $\mathbf{\Lambda}^{\frac{1}{2}}_{S\times S}$, $\mathbf{V}_{p\times S}$.

2. Calculate the matrix of residuals $\hat{\boldsymbol{\varepsilon}}_{n\times p} = \mathbf{X} - \hat{\mathbf{X}}^{(S)}$, and estimate σ^2.

3. For $b = 1, .., B$:

   (a) Draw ε_{ij}^b from $\mathcal{N}(0, \hat{\sigma}^2)$ to obtain a new matrix $\boldsymbol{\varepsilon}^b$,

   (b) Generate a new data table: $\mathbf{X}^b = \hat{\mathbf{X}}^{(S)} + \boldsymbol{\varepsilon}^b$, and

   (c) Perform PCA on $\mathbf{X}^b$ to obtain new estimates for the parameters $(\mathbf{U}^b, (\boldsymbol{\Lambda}^b)^{1/2}, \mathbf{V}^b)$ and a new estimator $\hat{\mathbf{X}}^{(S)^b} = \mathbf{U}^b(\boldsymbol{\Lambda}^b)^{\frac{1}{2}}\mathbf{V}^{b\prime}$.

This algorithm requires an estimate for the noise variance σ^2 in step 2. Since the maximum likelihood estimator is biased, Josse and Husson (2011) suggested the use of

$$\hat{\sigma}^2 = \frac{\|\mathbf{X} - \hat{\mathbf{X}}\|^2}{np - nS - pS + S + S^2},$$

which corresponds to the sum of squares of the residuals divided by the number of observed values minus the number of independent parameters estimated. At the end of the parametric bootstrap procedure, we have B sets of estimates for the parameters $((\mathbf{U}^1, (\boldsymbol{\Lambda}^1)^{1/2}, \mathbf{V}^1), ..., (\mathbf{U}^B, (\boldsymbol{\Lambda}^{1/2})^B, \mathbf{V}^B))$ and B fitted matrices $(\hat{\mathbf{X}}^{(S)^1} = \mathbf{U}^1(\boldsymbol{\Lambda}^{\mathbf{1}})^{1/2}\mathbf{V}^{1\prime}, ..., \hat{\mathbf{X}}^{(S)^B} = \mathbf{U}^B(\boldsymbol{\Lambda}^{\boldsymbol{B}})^{1/2}\mathbf{V}^{B\prime})$. We will discuss in Section 5 how to translate them into practical confidence areas.

Since as shown in Section 3 the PCA estimator is only asymptotically unbiased ($\mathbb{E}(\hat{\mathbf{X}}^{(S)}) = \tilde{\mathbf{X}}$), we expect the coverage of the parametric bootstrap confidence areas to reach their nominal level in an asymptotic framework. Huet et al (1999) studied properties of confidence intervals for the parameters of non-linear models in the asymptotic framework where the variance of the errors tends to zero, and showed that the level of confidence intervals based on the parametric residuals bootstrap are identical up to the order $o(\sigma)$ to the level of intervals calculated with first order asymptotics. Thus, in a low noise regime, we expect the parametric bootstrap procedure to give similar results to asymptotic method described in Section 3.

4.2 A Cell-wise Jackknife

The jackknife is another non-parametric way to study the variability of parameters. In the context of PCA, the "classical jackknife" procedure involves deleting each row of the data matrix $\mathbf{X}$ one at a time. Daudin et al (1989) and Besse and de Falguerolles (1993) used such a procedure to define

a stability index and select the number of PCA dimensions S. To cut the computational costs, Daudin et al (1989) used the infinitesimal jackknife (IJ) method whereas Besse and de Falguerolles (1993) used perturbation theory to define approximations.

Removing one row at a time is analogous to the non-parametric bootstrap where the rows are considered as a random sample. Here, in the setup where the rows are not *i.i.d*, we propose a new form of jackknife for PCA, which consists in removing one cell x_{ij} of the data matrix at a time and estimating the PCA parameters from the incomplete data set. Writing $\hat{\mathbf{X}}^{(S)^{(-ij)}}$ for the PCA estimator obtained from the matrix without the cell (ij), we get the pseudo-values

$$\hat{\mathbf{X}}_{jackk}^{(S)^{(ij)}} = \hat{\mathbf{X}} + \sqrt{np}\left(\hat{\mathbf{X}}^{(S)^{(-ij)}} - \hat{\mathbf{X}}\right). \tag{6}$$

This procedure is repeated for each cell of the data matrix. These pseudo-values can then be transformed into confidence ellipsoids as shown in Section 5.

Our jackknife procedure requires a method for performing PCA with missing values. A common approach to dealing with missing values in PCA consists in minimizing the loss function over all non missing entries:

$$\mathcal{C} = \sum_{i=1}^{n}\sum_{j=1}^{p} w_{ij}(x_{ij} - \sum_{s=1}^{S}\sqrt{d_s}q_{is}r_{js})^2,$$

where $w_{ij} = 0$ if x_{ij} is missing and $w_{ij} = 1$ otherwise. Unlike in the complete case, there is no explicit solution to this optimization problem and it is necessary to use iterative algorithms such as alternating weighted least squares algorithms (Gabriel and Zamir, 1979) or the expectation-maximization PCA (EM-PCA) algorithm (Kiers, 1997; Josse and Husson, 2012). This latter consists in setting the missing elements at initial values, performing the PCA on the completed data set, filling-in the missing values using the fitted values ($\hat{x}_{ij}^{(S)} = \sum_{s=1}^{S}\sqrt{\lambda_s}u_{is}v_{js}$) using a predefined number of dimensions S and repeating the procedure on the newly obtained matrix until the total change in the matrix falls below an empirically determined threshold.

Formally, the jackknife tells us about the variability of the mean of the leave-one-out estimates. In most applications, the mean of the leave-one-out estimates is very close to the actual estimate (in the case of the sample mean, they are the same). Consequently, this specific jackknife is a competitor to both the bootstrap procedure and the method based on the asymptotic results in the sense that they estimate the same variability. Efron and Stein (1981, Theorem 2) showed that

the jackknife estimate of variance is biased upwards in expectation, suggesting that the jackknife should lead to conservative confidence areas.

4.3 Approximating the Jackknife

The jackknife procedure described above is computationally demanding, as it effectively requires us to run the iterative EM-PCA algorithm $n \times p$ times. To reduce the computational effort required, we consider an approximation to the jackknife based on the "leave-out-one" lemma of Craven and Wahba (1979).

Josse and Husson (2011) showed that the prediction error $(x_{ij} - \hat{x}_{ij}^{(S)^{(-ij)}})$ for a cell (ij) can be approximated by

$$x_{ij} - \hat{x}_{ij}^{(S)^{(-ij)}} \simeq \frac{x_{ij} - \hat{x}_{ij}}{1 - \mathbf{P}_{ij,ij}}. \tag{7}$$

Thus, we can define an approximation to the jackknife where instead of deleting the cell (ij), we replace it with the estimate $\hat{x}_{ij}^{(S)^{(-ij)}}$ obtained using (7) and then perform PCA on this altered matrix. The pseudo-values are computed as in (6). The advantage of this approach is that it allows us to avoid solving missing value PCA problems.

Note that performing PCA on the data matrix where the cell (ij) is missing (using the EM-PCA algorithm for instance) is equivalent to performing PCA on a data set where the cell (ij) has been replaced with $\hat{x}_{ij}^{(-ij)}$. Consequently, if the approximation (7) were exact, the procedure outlined here would be equivalent to the jackknife.

5 Visualizing Confidence Areas

All our proposed methods effectively provide a list of pseudo-realizations of the PCA estimator

$$\left\{\hat{\mathbf{X}}^{(S)^1} = \mathbf{U}^1(\mathbf{\Lambda^1})^{1/2}\mathbf{V}^{1\prime}, ..., \hat{\mathbf{X}}^{(S)^\star} = \mathbf{U}^\star(\mathbf{\Lambda^\star})^{1/2}\mathbf{V}^{\star\prime}\right\}, \tag{8}$$

where $\star = B$ for the parametric bootstrap and $\star = np$ for the jackknife. For example, with the bootstrap, each pseudo-realization $\hat{\mathbf{X}}^{(S)^b}$ corresponds to the rank-S approximation obtained by performing PCA on the noised matrix $\mathbf{X}^b = \hat{\mathbf{X}}^{(S)} + \boldsymbol{\varepsilon}^b$. The spread of these pseudo-realizations can then give us an idea of the stability of PCA.

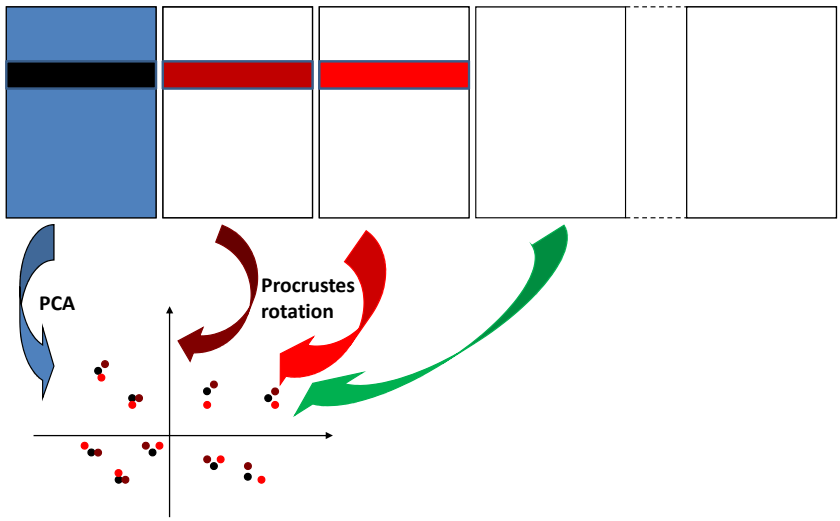

Figure 1: Procrustes rotations of the rows of the PCA estimators onto the initial configuration. The first table corresponds to $\mathbf{X}$ and the other to $\hat{\mathbf{X}}^{(S)^1}, ..., \hat{\mathbf{X}}^{(S)^\star}$.

In order to make good use of the pseudo-realizations, we need to find a way to combine them into confidence areas and visualize the results. As mentioned in the introduction, this is not a trivial task since the principal components can be rotated with respect to each other from one configuration to another. At a high level, if we want to visualize confidence intervals for the rows (or columns), we need to transform $\hat{\mathbf{X}}^{(S)^b}$ into an approximation $\breve{\mathbf{X}}^{(S)^b}$ whose rows (columns) are contained in the row- (column-) span of $\hat{\mathbf{X}}^{(S)}$. Once we have made this transformation, we can generate confidence ellipses by looking at the variance of the $\breve{\mathbf{X}}^{(S)^b}$.

We address the problem of transforming $\hat{\mathbf{X}}^{(S)^b}$ into the appropriate space of $\hat{\mathbf{X}}^{(S)}$ using Procrustes rotation. (Gower and Dijksterhuis, 2004; Krzanowski, 2010; Schoenemann, 1966), which is often used to compare the results of different PCA analyses (Lebart, 2007). Recall that all our matrices are already centered, so we do not need to worry about translation. Suppose that we want to study the rows of $\mathbf{X}$. Then, we define $\breve{\mathbf{X}}^{(S)^b}$ as

$$\breve{\mathbf{X}}^{(S)^b} = \hat{\mathbf{X}}^{(S)^b}\mathbf{R}^b \text{ where } \mathbf{R}^b = \text{argmin}_{\mathbf{R}}\left\{\left\|\hat{\mathbf{X}}^{(S)} - \hat{\mathbf{X}}^{(S)^b}\mathbf{R}\right\|^2 : \mathbf{R}'\mathbf{R} = \mathbb{I}_p\right\}, \tag{9}$$

i.e., $\breve{\mathbf{X}}^{(S)^b}$ is the best rotation of $\hat{\mathbf{X}}^{(S)^b}$ towards $\hat{\mathbf{X}}^{(S)}$. The well-know solution to (9) is to set $\mathbf{R}^b = \mathbf{E}_l\mathbf{E}_r'$, where $\mathbf{E}_l$ and $\mathbf{E}_r$ represent the left and right singular vectors of the matrix $(\hat{\mathbf{X}}^{(S)^b})'\hat{\mathbf{X}}^{(S)}$. As desired, the rows of $\breve{\mathbf{X}}^{(S)^b}$ are contained in the row-span of $\hat{\mathbf{X}}^{(S)}$ by construction.

Finally, we note that our procedure relies heavily on our ability to generate pseudo-realizations

of the form (8). Thus, it cannot be applied to the non-parametric bootstrap (which involves resampling rows) or to the classical jackknife (which deletes one row at a time). Both of these approaches modify the structure of the rows of $\mathbf{X}$ for each bootstrap replication, and so do not produce quantities of the form (8).

6 Results

To assess our proposed methods, we first run a comparative simulation study below, and then apply the methods to a wine dataset in Section 6.2 and a decathlon dataset in Section 6.3.

6.1 Simulations

For our simulation study, we generated data according to the model (2) with $S = 2$ dimensions. More precisely, $\tilde{\mathbf{X}}$ is generated as follows:

1. A SVD is performed on an $n \times S$ matrix generated from a standard multivariate normal distribution. The left singular vectors provide S empirically orthonormal vectors.

2. Each vector $s : 1, ..., S$ is replicated to obtain p variables such that the ratio between the eigenvalues is fixed to a pre-determined value d_1/d_2. For instance, if $p = 50$ and $d_1/d_2 = 4$, the first vector is replicated 40 times and the second vector is replicated 10 times.

3. Then, to generate the data matrix $\mathbf{X}$, we add an isotropic Gaussian noise to the structure. The scale of the Gaussian noise is determined by the target signal-to-noise (SNR) ratio.

Here, a high SNR implies that the columns of $\mathbf{X}$ are highly correlated, whereas a low SNR implies that the data is very noisy. For each combination of the parameters, we generated 50 data sets. Having generated data matrices $\mathbf{X}$, we ran all our proposed methods to generate confidence ellipses for the PCA representation as discussed in Section 5, and then evaluated the coverage of the ellipses. To do so, we counted the number of times that the row points of the true configuration $\tilde{\mathbf{X}}$ are inside the ellipses. Results of the simulation study are gathered in Table 1.

As our simulation study makes clear, the proposed methods have very different strengths and weaknesses. The first two methods, namely the asymptotic method and the parametric bootstrap, both perform similarly although the bootstrap is somewhat better. Both methods provide accurate

Table 1: Simulated coverage for the asymptotic method, the bootstrap, the jackknife, and the approximation of the jackknife. Nominal coverage is $1 - \alpha = 0.95$, and results are averaged over 50 replicates. We vary the signal-to-noise ratio (SNR), the relative scale of the first and second principal components (d_1/d_2), and the number of variables p and rows n. In bold, values greater than 0.93 and less than 0.97; in italic values less than 0.90. All simulations use a two-dimensional Gaussian signal as the true signal. In some cases, the exact jackknife was too slow so we only ran its approximation.

SNR	d1/d2	p	n	Asymp	Boot	Jack	Approx Jack
1.00	1.00	5	20	*0.889*	0.92	0.919	*0.863*
1.00	1.00	5	50	0.916	0.928	*0.836*	*0.817*
1.00	1.00	5	100	**0.931**	**0.935**	*0.797*	*0.788*
1.00	1.00	20	20	*0.851*	0.902	**0.948**	0.924
1.00	1.00	20	50	0.917	**0.935**	0.925	0.916
1.00	1.00	20	100	**0.93**	**0.94**	-	0.917
1.00	1.00	100	20	*0.786*	*0.852*	-	0.97
1.00	4.00	5	20	*0.856*	*0.892*	**0.962**	*0.896*
1.00	4.00	5	50	0.913	0.927	**0.959**	0.914
1.00	4.00	5	100	**0.931**	**0.937**	**0.956**	0.923
1.00	4.00	20	20	*0.81*	*0.892*	**0.957**	**0.931**
1.00	4.00	20	50	0.901	**0.931**	0.919	0.905
1.00	4.00	20	100	0.924	**0.936**	-	*0.896*
1.00	4.00	100	20	*0.716*	*0.835*	-	0.976
4.00	1.00	5	20	**0.938**	**0.938**	*0.868*	*0.849*
4.00	1.00	5	50	**0.935**	**0.935**	*0.805*	*0.796*
4.00	1.00	5	100	**0.94**	**0.941**	*0.769*	*0.765*
4.00	1.00	20	20	**0.938**	**0.941**	**0.951**	**0.933**
4.00	1.00	20	50	**0.943**	**0.944**	**0.933**	0.927
4.00	1.00	20	100	**0.942**	**0.945**	-	0.925
4.00	1.00	100	20	**0.935**	**0.936**	-	**0.953**
4.00	4.00	5	20	0.926	0.924	1	0.999
4.00	4.00	5	50	**0.943**	**0.941**	0.998	0.996
4.00	4.00	5	100	**0.938**	**0.937**	0.999	0.998
4.00	4.00	20	20	**0.934**	**0.938**	**0.945**	0.929
4.00	4.00	20	50	**0.938**	**0.941**	0.91	0.902
4.00	4.00	20	100	**0.944**	**0.943**	-	0.9
4.00	4.00	100	20	0.925	**0.932**	-	**0.957**

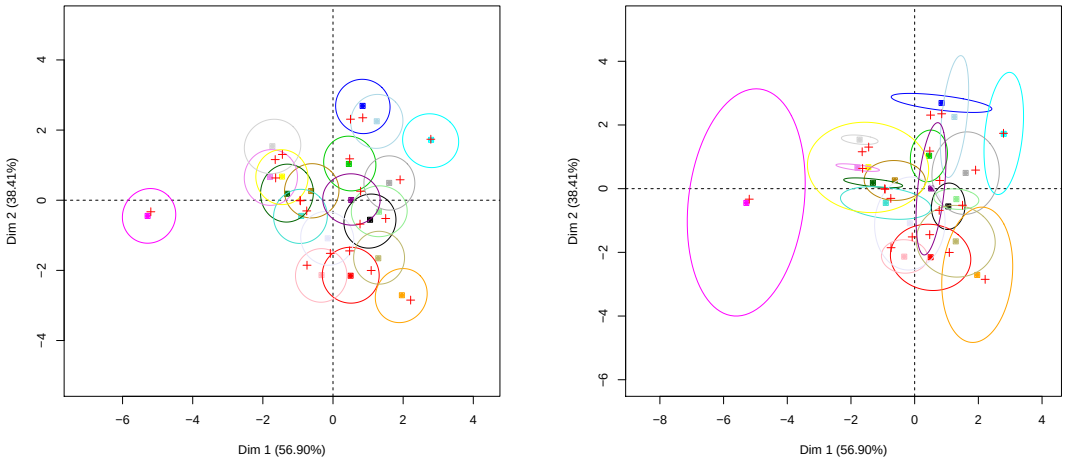

Parametric Bootstrap Jackknife

Figure 2: Simulation with $n = 20$, $p = 5$, $SNR = 4$, $d1/d2 = 1$. The ellipses represent confidence areas for the two-dimensional representations of the row points produced by PCA. The row points of the true configuration $\tilde{\mathbf{X}}$ are represented with red crosses.

confidence areas close to the nominal level when the SNR is high and when n is large, but can break down in low SNR settings. This is not surprising, as the asymptotic method was justified by a linear approximation which is valid in the low noise limit, and the parametric bootstrap should behave similarly to the asymptotic method when both are valid.

Conversely, the jackknife and its approximation behave well in low SNR settings, but can struggle when p is small or when there is not much noise. In general, the exact jackknife is more accurate, but can sometimes be prohibitively expensive computationally.

To gain insight into these results we focus on specific cases. First, we consider a small matrix ($n = 20$, $p = 5$, SNR=4, $d_1/d_2 = 1$) for which the jackknife had a lot of trouble. As seen in Figure 2, the confidence ellipsoids produced by the jackknife are highly variable: some of them are much too big, whereas others are much too small. This variability then makes the jackknife have poor average coverage. Meanwhile, the bootstrap here achieves stable variance estimates and good coverage.

The reason for this instability appears to be that the jackknife ellipses are each dominated by very few pseudo-realizations. It seems that the ellipse for each row point is governed by only $p = 5$ pseudo-realizations obtained by leaving out cells from that row. Moreover, even the few pseudo-realizations we have appear to all be fairly correlated. This may be an artifact of our

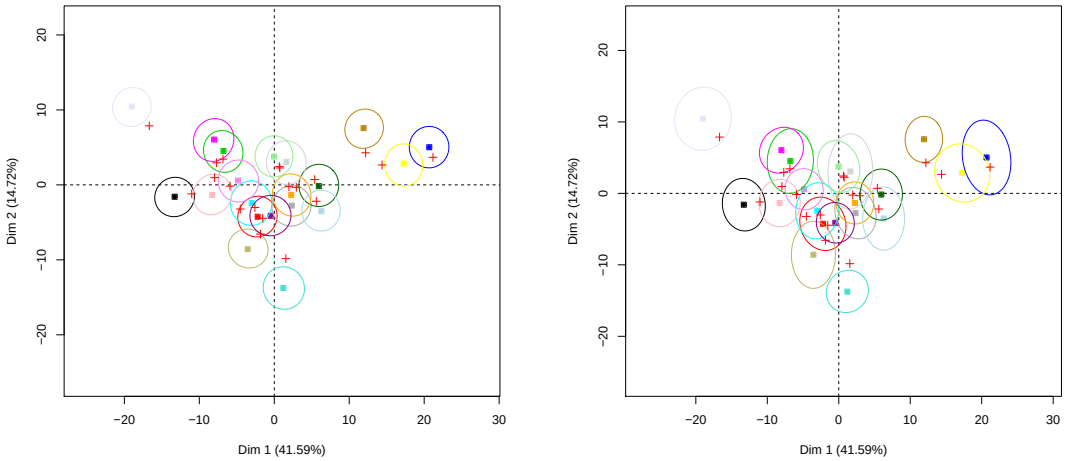

Parametric Bootstrap Jackknife

Figure 3: One simulation with $n = 20$, $p = 100$, $SNR = 1$, $d1/d2 = 4$. The ellipses represent confidence areas for the two-dimensional representations of the row points produced by PCA. The row points of the true configuration $\tilde{\mathbf{X}}$ are represented with red crosses.

data generation procedure, where many columns are almost identical in the low noise setup. This correlation between columns could also explain why the jackknife deals rather poorly with high SNR in our simulation, even when p is larger than 5.

Figure 3 illustrates a lower signal case ($n = 20$, $p = 100$, SNR=1, $d_1/d_2 = 4$), for which the bootstrap has more trouble. We see that, here, the bootstrap appears to be over-fitting the data and under-estimating the size of the noise; the bootstrap ellipses are too small and often do not contain the true values. Meanwhile, the jackknife appears to be performing reasonably well here. In repeated simulations (not shown in Table 1), the methods achieved average coverage proportions of 0.79 for the asymptotic method, 0.83 for the bootstrap, 0.91 for the jackknife and 0.89 for the approximate jackknife.

6.2 Wine Data

We next apply the proposed methods on a dataset from sensory analysis on wines. The data is available from `http://factominer.free.fr/docs/Consumer_wine.csv`. The data was collected by asking consumers to describe 10 white wines (5 Sauvignon and 5 Vouvray) with 15 sensory attributes including sweetness and astringency: each consumer gave each wine a score between 1

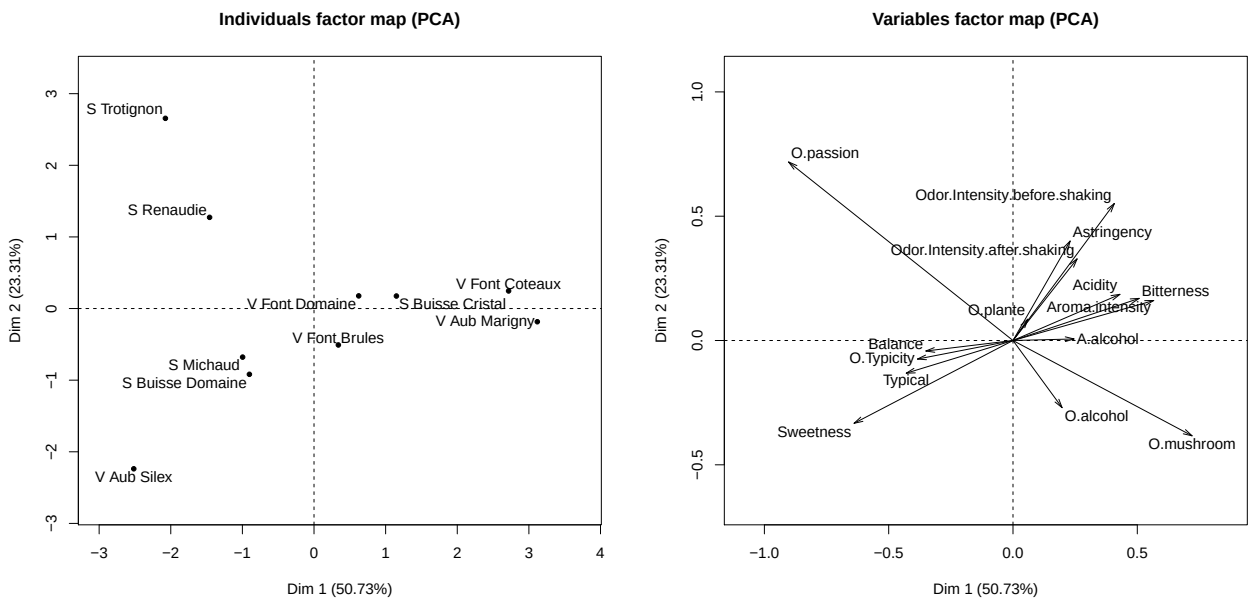

(a) 2 dimensional wines representation.

(b) PCA variables.

Figure 4: PCA on the wine dataset.

and 10 for each attribute. The data was then averaged across consumers into a 10×15 data matrix $\mathbf{X}$, for which we want to perform PCA. Figure 4 shows two-dimensional PCA representations, the rows coordinates ($\mathbf{U\Lambda^{(1/2)}}$) as well as variables coordinates ($\mathbf{V\Lambda^{(1/2)}}$) for this dataset.

Our goal is to associate confidence areas with the low-rank representations shown in Figure 4. Here the 10 wines and 15 attributes are all non-random, so methods based on re-sampling rows of $\mathbf{X}$ would not be applicable. Moreover, the effective noise in $\mathbf{X}$ gets smaller as the number of interviewed consumers grows; thus the low-noise asymptotic regime studied in Section 3 makes sense for this example.

Figure 5 show the result of applying our proposed methods to this dataset. Note that these methods all require us to specify the number of dimensions S as an input parameter. Different methods for doing this are available in the literature (Jolliffe, 2002); here, we used a cross-validation procedure (Josse and Husson, 2011) which suggested that using $S = 2$ dimensions was appropriate.

The confidence ellipses produced by the asymptotic method and by the bootstrap are all fairly round. In comparison, the ellipses obtained by the jackknife are larger and are much more elongated. In particular, the jackknife appears to suggest that, for the individuals (Figure 5),

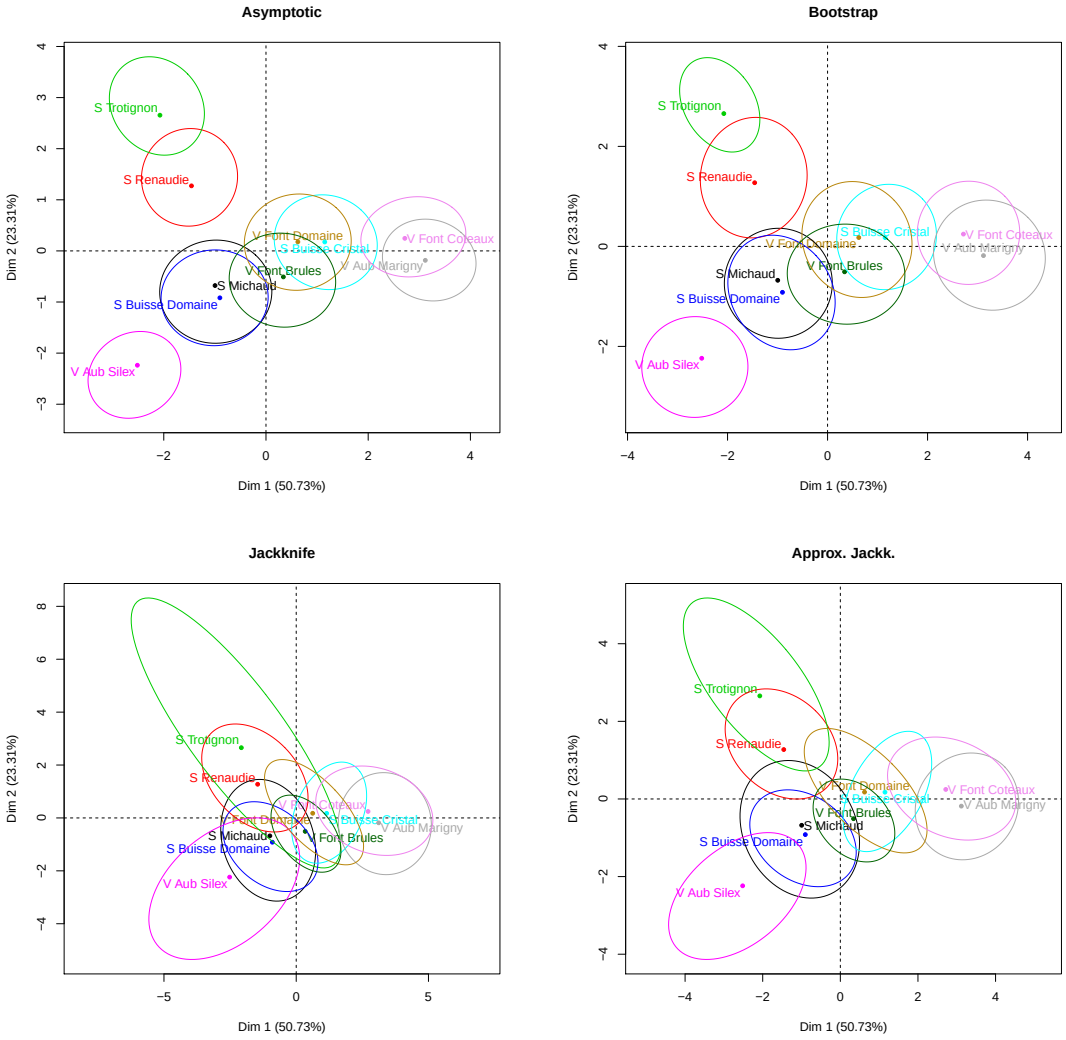

Figure 5: Confidence areas around the wines of the PCA representation using the asymptotic variance, the parametric bootstrap, the jackknife and the approximation of the jackknife.

Table 2: Simulation study, using the rank-two approximation to the wine data matrix as the true signal. The table shows average coverage for the produced confidence ellipsoids; nominal coverage is $1 - \alpha = 0.95$. The true matrix was perturbed with isotropic Gaussian noise of various scales σ. These results are averaged over 200 replicates.

Noise scale σ	Asymptotic	Bootstrap	Jackknife	Approx. Jack.
0.1	**0.942**	**0.938**	0.974	**0.95**
0.2	**0.938**	**0.938**	0.98	**0.952**
0.4	0.916	**0.93**	0.984	**0.958**
0.6	*0.876*	0.92	0.982	**0.953**
0.8	*0.804*	*0.89*	0.986	**0.948**
1	*0.726*	*0.838*	0.982	**0.936**
1.2	*0.64*	*0.816*	0.986	**0.934**

there is much less signal in the upper-left direction dominated by the Sauvignon Trotignon and the Sauvignon Renaudie wines than in the other directions. Conversely, all methods seem to imply that the Vouvray Aubence Silex wine in the lower-left corner is perceived differently from most other wines in a significant way, especially from the point of view of the descriptor sweetness.

In order to get a better idea of which method was most trustworthy here, we ran a small simulation study built on top of the wine dataset. We used the fitted rank-two PCA matrix as the true signal matrix, and then added isotropic Gaussian noise to it. Results for various noise scales σ are shown in Table 2. The root-mean residual for the actual PCA fit to the wine data was $\sigma = 0.6$. Our small simulation seems to suggests that the bootstrap was slightly anti-conservative here, whereas the jackknife was slightly conservative. Note the good behavior of the approximate jackknife. Thus, the most honest confidence areas for the wine dataset should probably be somewhere in between the bootstrap and jackknife ellipses shown in Figure 5.

6.3 Decathlon Data

So far, we have focused on the case where the number of dimension S is equal to 2. Indeed, when PCA is used as a visualization tool, it is very common to examine and interpret only the first two dimensions. This can be explained in part because it is easier from a visualization point of view but also because it makes sense in many applications. However some datasets cannot be summarized in two dimensions. For instance, let us consider the decathlon dataset (available in the R package FactoMineR (Lê et al, 2008)) which contains the performance of 41 athletes in the

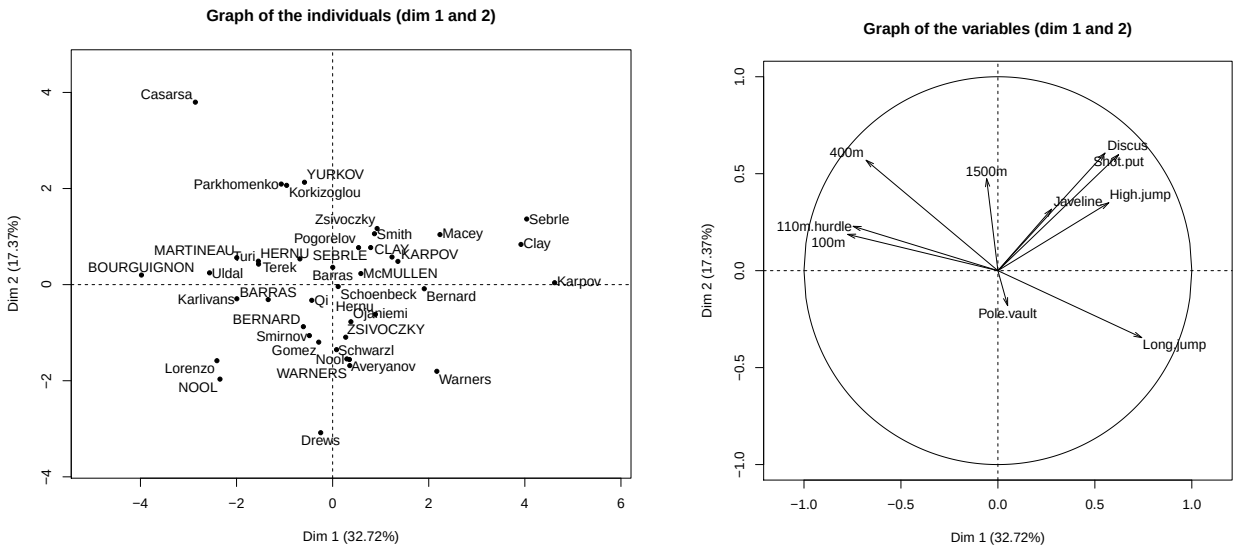

Figure 6: Decathlon dataset: representation of the individuals (left) and of the variables (right) on the first two dimensions of the PCA.

different disciplines of the decathlon during the Olympic Game 2004 and the Decastar 2004. The first dimension (see Figure 6) rank the athletes from the best ones on the right who performed well on many trials (they run fast, jump high, threw the discuss far, etc.) to the ones with lower results on the left. The second dimension opposes the powerful athletes (high values for discuss and shot.put) to the others. The first two dimensions represent 50% of the variability. A greater percentage is not desirable since it would imply that the decathlon does not require different physical aptitudes (if all the trials were highly correlated). The dimensions 3 and 4 (see Figure 7) allow to differentiate athletes who score high at the 1500m and pole vault against those who shine in the javelin trial.

The proposed methods (Section 4) can be applied straightforwardly to study the variability of the low rank estimator even for a rank greater than 2. We illustrate this ability using a small simulation study built on the top of the decathlon dataset in the same way as the simulations for the wine dataset (Section 6.2): an isotropic Gaussian noise is added to the signal in 4 dimensions. Using all the pseudo-realizations, we built ellipsoids in 4 dimensions and we counted the number of times that the row points of the signal are inside the ellipsoids to assess the coverage. The results obtained in Table 3 are in agreement with the previous results showing results close to the nominal level for all the methods with a slightly underestimation for the bootstrap and an

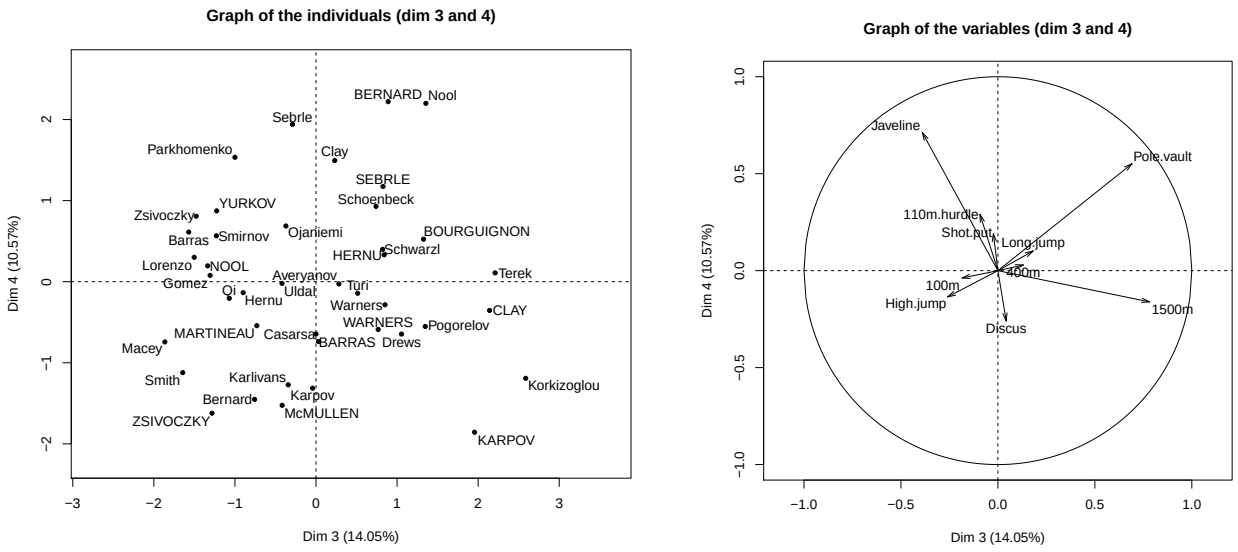

Figure 7: Decathlon dataset: representation of the individuals (left) and of the variables (right) on the dimensions 3 and 4 of the PCA.

overestimation for the jackknife in high level noise situations. The approximate jackknife shows a good behavior in all the simulations.

Then we visualize the results on the real dataset representing the ellipsoids around the individuals on the PCA representation on the dimensions 1-2 and on the dimensions 3-4 (Figures 8). We do not represent the asymptotic method since the results are very close to the ones obtained by the bootstrap. Figure 8 gives the projections of 4-dimensional ellipsoids in 2 dimensions and thus such areas can only be interpreted as areas of variability and not anymore as confidence areas. The ellipses obtained by all the methods are quite large suggesting a noisy data set. In such a noisy situation, the simulations (Table 3) indicated the bad behavior of the Jackknife with a coverage close to one. This is in agreement with what is observed where some individuals have very large ellipses which leads to graphics that are impossible to interpret and which seem difficult to trust. There are a lot of overlapping between the ellipses on the dimensions 3-4 for the bootstrap and the approximate Jackknife which may suggest that these dimensions are not very significant. Nevertheless, these ellipses allow to differentiate between individuals such as Karpov on the bottom to Nool on the top.

Table 3: Simulation study, using the rank-four approximation to the decathlon data matrix as the true signal. The table shows average coverage for the produced confidence ellipsoids; nominal coverage is $1 - \alpha = 0.95$. The true matrix was perturbed with isotropic Gaussian noise of various scales σ. These results are averaged over 200 replicates.

Noise scale σ	Asymptotic	Bootstrap	Jackknife	Approx. Jack.
0.05	**0.946**	**0.941**	**0.942**	**0.934**
0.1	**0.941**	**0.941**	**0.945**	**0.935**
0.2	**0.936**	**0.939**	**0.949**	**0.941**
0.3	**0.937**	**0.939**	**0.952**	**0.942**
0.4	**0.932**	**0.938**	**0.957**	**0.941**
0.6	0.91	0.921	**0.963**	**0.939**
0.8	*0.897*	0.912	0.973	**0.946**
1	*0.86*	*0.891*	0.98	**0.935**
1.2	*0.827*	*0.862*	0.982	**0.931**

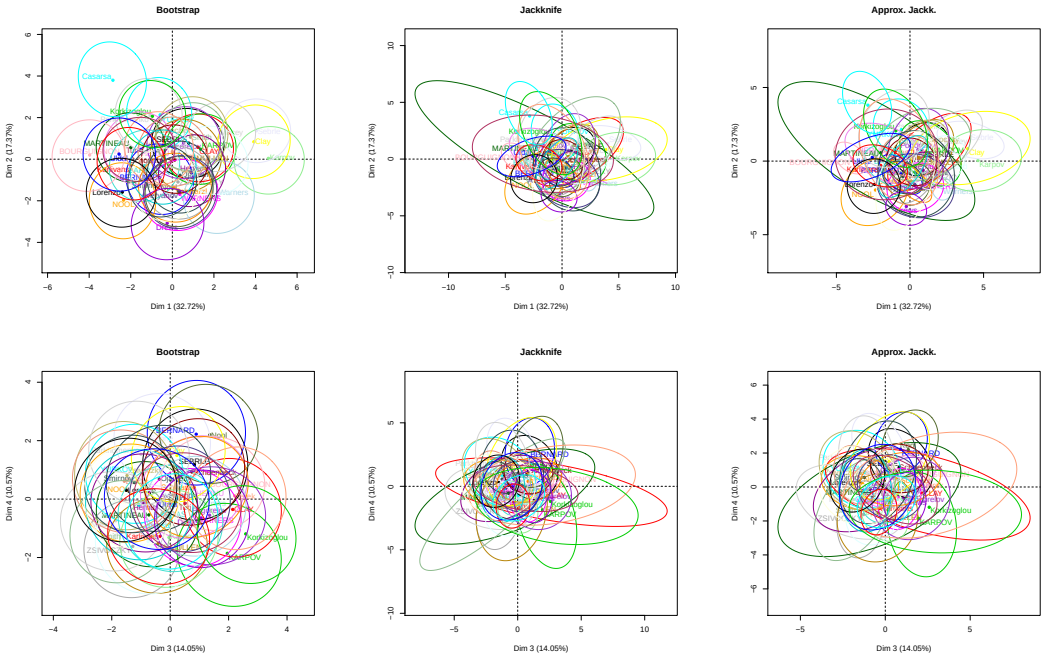

Figure 8: Confidence areas around the athletes of the PCA representation using the parametric bootstrap, the jackknife and the approximation of the jackknife for the dimensions 1-2 on the top and 3-4 on the bottom.

7 Discussion

In this paper, we introduced and studied various approaches to estimating the variance of the PCA estimator and suggested a solution for visualizing these estimates as confidence ellipsoids on PCA maps, emphasizing that confidence areas ensure relevant interpretation. We compared coverage properties of confidence areas based on an asymptotic method, a parametric bootstrap, a jackknife and an approximation to the jackknife, and found that the resulting variance estimates vary widely. The asymptotic method and the bootstrap do well in low-noise setting, but can fail when the noise level gets high or when the number of variables is much greater than the number of rows. On the other hand, the jackknife has good coverage properties for large noisy examples but requires a minimum number of variables to be stable enough.

All our methods rely on the validity of the model (2), and consequently the choice of the number of dimension S is crucial. We focused mainly on the case where $S = 2$ on the simulations. However, we showed that the methods can still be applied and are accurate for a number greater than 2. The results are then visualized with areas of variability.

Note that the principal components of $\mathbf{X}$ obtained by PCA are unbiased estimates of the true principal components of $\tilde{\mathbf{X}}$ in the asymptotic regime when the noise variance tends to 0. But when the signal to noise ratio is small, the principal components of $\mathbf{X}$ are biased and consequently the confidence regions centered on these principal components are not necessarily optimal. Nevertheless, even when they can no longer be considered as confidence regions for the true underlying principal components, they can still be used to visualize the variance of the PCA estimator due to noise.

There appear to be many potential applications for our methods. For example, in the field of plant breeding, biplots are often used to represent the interaction between genotypes and environments; in this case, having access to confidence areas on the biplot representation is important for protection against spurious interpretations about, say, differences between genotypes. Other approaches to this problem include the work of Perez-Elizalde et al (2011) and Josse et al (2014) who suggested a Bayesian treatment of the model enabling them to get direct distributions for the parameters, and that of Yang et al (2013) who proposed another bootstrap approach which consists in simultaneously bootstrapping rows and columns of the data matrix.

We finish by discussing some opportunities for further research. Extending our method to

account for uncertainty in the number of dimensions S presents interesting challenges. One idea would be be to build on ideas of Efron (2013), who discusses the estimation of prediction error when the model that was fit is itself a random variable. Another solution would be to use a fully Bayesian approach in the line with the works of Hoff (2007, 2009) for models based on singular-value decompositions. Visualizing confidence areas when S is random may also prove to be difficult

Second, our methods could be used for studying the variability of other estimators related to PCA. For instance, several estimators that apply different thresholding and/or shrinking rules on the empirical singular values while keeping the empirical singular vectors unchanged have been proposed for denoising data in the context of low-rank matrix estimation (Mazumder et al, 2010; Verbanck et al, 2013; Donoho and Gavish, 2013; Chatterjee, 2013; Josse and Sardy, 2014; Rao., 2013; Shabalin and Nobel, 2013). Shabalin and Nobel (2013) proposed an estimator that explicitly accounts for the effect of the noise on the singular values and on the singular vectors. This family of estimators can provide a better recovery of the underlying signal in term of mean squared error in comparison to PCA. These estimators reduce variance at the price of bias. Studying the variability of these estimators with the proposed approaches as well as their coverage properties could lead to useful insights.

Finally, we note that it may be worth using the values of the projection matrix $\mathbf{P}_{ij,ij}$ as leverage values. Of course, these can only be seen as approximate measures of leverage since they are based on a linear approximation to the expectation surface. Nevertheless, they can be used to define corrected residuals $\hat{\varepsilon}_{ij} = \frac{\hat{x}_{ij} - x_{ij}}{\sqrt{1-\mathbf{P}_{ij,ij}}}$ or Cook's distances in the framework of PCA. In addition, in non-linear regression, such values are related to curvature measures (Laurent and Cool, 1993), and can be used to assess the impact of leverage points and outliers on the jackknife (Simonoff and Tsai, 1986). More research can be done in these directions.

Acknowledgements

Julie Josse has received the support of the European Union, in the framework of the Marie-Curie FP7 COFUND People Programme, through the award of an AgreenSkills fellowship (under grant agreement n 267196) for an academic visit to Stanford.

References

Badamoradi H, van den Berg F, Rinnan A (2013) Bootstrap based confidence limits in principal component analysis - a case study. Chemometrics and Intellligent Laboratory Systems 120:97–105

Bartholomew DJ (1987) Latent Variable Models and Factor Analysis. Griffin

Bates D, Watt D (1980) Relative curvature measure of nonlinearity. Journal of the Royal Society series B 42:1–25

Besse P, de Falguerolles A (1993) Application of resampling methods to the choice of dimension in principal component analysis. In: Computer Intensive Methods in Statistics, Physica-Verlag, pp 167–176

Candes E, Tao T (2009) The power of convex relaxation: Near-optimal matrix completion. IEEE Trans Inform Theory 56(5):2053–2080

Caussinus H (1986) Models and uses of principal component analysis (with discussion). In: de Leeuw J, Heiser W, Meulman J, Critchley F (eds) Multidimensional Data Analysis, DSWO Press, pp 149–178

Chateau F, Lebart L (1996) Assessing sample variability in the visualization techniques related to principal component analysis: bootstrap and alternative simulation methods. In: Prats A (ed) COMPSTAT, Physica-Verlag, pp 205–210

Chatterjee S (2013) Matrix estimation by universal singular value thresholding. arXiv:12121247

Cornelius P, Crossa J, Seyedsadr M (1996) Statistical tests and estimators of multiplicative models for genotype by environment interaction. In: Genotype by environment interaction, M.S Kang and H.G Gauch (eds), CRC press, Boca Raton, FL, pp 199–234

Craven P, Wahba G (1979) Smoothing noisy data with spline functions. Numer Math 31(4):377–403

Daudin JJ, Duby C, Trecourt P (1988) Stability of principal component analysis studied by the bootstrap method. Statistics: A journal of theoretical and applied statistics 19(2):241–258

Daudin JJ, Duby C, Trecourt P (1989) Pca stability studied by the bootstrap and the infinitesimal jackknife method. Statistics: A journal of theoretical and applied statistics 20(2):255–270

Denis JB, Gower JC (1994) Asymptotic covariances for the parameters of biadditive models. Utilitas Mathematica pp 193–205

Denis JB, Gower JC (1996) Asymptotic confidence regions for biadditive models: interpreting genotype-environment interactions. Applied Statistics 45(4):479–493

Denis JB, Pazman A (1999) Bias of least squares estimators in nonlinear regression models with constraints. part ii: Biadditive models. Applications of Mathematics 44:359–374

Donoho DL, Gavish M (2013) The optimal hard threshold for singular values is 4/sqrt(3). arXiv:13055870

Efron B (2013) Estimation and accuracy after model selection. Journal of the American Statistical Association (just-accepted)

Efron B, Stein C (1981) The jackknife estimate of variance. Annals of Statistics 3:586–596

Efron B, Tibshirani R (1994) An Introduction to the Bootstrap. Chapman & Hall/CRC

Gabriel KR, Zamir S (1979) Lower rank approximation of matrices by least squares with any choice of weights. Technometrics 21(4):236–246

Gauch H (1988) Model selection and validation for yield trials with interaction. Biometrics 44:705–715

Gauch H, Zobel R (1996) Ammi analysis of yield trials. In: Genotype by environment interaction, M.S Kang and H.G Gauch (eds), CRC press, Boca Raton, FL, pp 141–150

Gower JC, Dijksterhuis GB (2004) Procrustes Problems. New York: Oxford University Press

Greenacre M (1984) Theory and Applications of Correspondence Analysis. Acadamic Press

Hoff PD (2007) Model averaging and dimension selection for the singular value decomposition. Journal of the American Statistical Association 102(478):674–685

Hoff PD (2009) Simulation of the matrix Bingham–von Mises–Fisher distribution, with applications to multivariate and relational data. Journal of Computational and Graphical Statistics 18(2):438–456

Holmes S (1989) Using the bootstrap and the rv coefficient in the multivariate context. In: Proceedings of the conference on Data analysis, learning symbolic and numeric knowledge, vol 24, pp 119–131

Huet S, Denis JB, Adamczyk K (1999) Bootstrap confidence intervals in nonlinear regression models when the number of observations is fixed and the variance tends to 0. application to biadditive model. Statistics 32:203–227

Jolliffe IT (2002) Principal Component Analysis. Springer

Josse J, Husson F (2011) Selecting the number of components in pca using cross-validation approximations. Computational Statististics and Data Analysis 56(6):1869–1879

Josse J, Husson F (2012) Handling missing values in exploratory multivariate data analysis methods. Journal de la Société Française de Statistique 153(2):79–99

Josse J, Sardy S (2014) Selecting thresholding and shrinking parameters with generalized SURE for low rank matrix estimation. arXiv 13106602

Josse J, van Eeuwijk F, Piepho HP, Denis JB (2014) Another look at bayesian analysis of ammi models for genotype-environment data. Journal of Agricultural, Biological, and Environmental Statistics pp 1–18

Kiers HAL (1997) Weighted least squares fitting using ordinary least squares algorithms. Psychometrika 62(2):251–266

Krzanowski WJ (2010) Principles of multivariate analysis; a user's perspective. Clarendon Press, Oxford

Laurent R, Cool R (1993) Leverage, local influence and curvature in nonlinear regression. Biometrika 80 (1):99–106

Lê S, Josse J, Husson F (2008) Factominer: An r package for multivariate analysis. Journal of Statistical Software 25(1):1–18, URL http://www.jstatsoft.org/v25/i01

Lebart L (2007) Which bootstrap for principal axes methods? In: Selected Contributions in Data Analysis and Classification, Springer Berlin Heidelberg, pp 581–588

Mazumder R, Hastie T, Tibshirani R (2010) Spectral regularization algorithms for learning large incomplete matrices. Journal of Machine Learning Research 99:2287–2322

Milan M (1995) Application of the parametric bootstrap to models that incorporate a singular value decomposition. Journal of the Royal Statistical Society Series C 44:31–49

Papadopoulo T, Lourakis MIA (2000) Estimating the jacobian of the singular value decomposition: Theory and applications. In: In Proceedings of the European Conference on Computer Vision, ECCV00, Springer, pp 554–570

Pazman A (2002) Results on nonlinear least squares estimators under nonlinear equality constraints. Journal of statistical planning and inference 103:401–420

Pazman A, Denis JB (2002) Measures of nonlinearity for biadditive anova models. Metrika 55 (3):233–245

Perez-Elizalde S, Jarquin D, Crossa J (2011) A general bayesian estimation method of linear-bilinear models applied to plant breeding trials with genotype x environment interaction. Journal of Agricultural Biological and Environmental Statistics 17(1):15–37

Rao NR (2013) Optshrink - low-rank signal matrix denoising via optimal, data- driven singular value shrinkage. arXiv:13066042

Schoenemann P (1966) A generalized solution of the orthogonal procrustes problem. Psychometrika 31(1):1–10

Shabalin AA, Nobel B (2013) Reconstruction of a low-rank matrix in the presence of Gaussian noise. Journal of Multivariate Analysis 118(0):67 – 76

Simonoff J, Tsai CL (1986) Jackknife-based estimators and confidence regions in nonlinear regression. Technometrics 28 (2):103–112

Timmerman ME, Kiers HAL, Smilde AK (2007) Estimating confidence intervals for principal component loadings: a comparison between the bootstrap and asymptotic results. British Journal of Mathematica and Statistical Psychology 60(2):295–314

Tipping M, Bishop CM (1999) Probabilistic principal component analysis. Journal of the Royal Statistical Society B 61(3):611–622

Verbanck M, Husson F, Josse J (2013) Regularized PCA to denoise and visualize data. Statistics and Computing forthcoming

Yang RC, Crossa J, Cornelius P, Burguenoc J (2013) Improved statistical inference for graphical description and interpretation of genotype environment interaction. Crops Sciences 53 (6):2400–2410